

1 CLAIMS

2
3 1. A system comprising:

4 a forwarding component that forwards packets;

5 a classifying component that classifies packets and is capable of classifying
6 packets for the forwarding component;

7 a session tracking component that tracks sessions for at least one of the
8 forwarding component and the classifying component;

9 a health and load handling component that is capable of providing health
10 and load information to the classifying component; and

11 a high availability mechanism that provides detection of, handling of, and
12 recovery from a failure of one or more of the forwarding component, the
13 classifying component, the session tracking component, and the health and load
14 handling component.

15
16 2. The system as recited in claim 1, wherein the system further
17 comprises:

18 a request routing component that is capable of routing logical requests;

19 wherein the high availability mechanism provides detection of, handling of,
20 and recovery from a failure of the request routing component.

1 3. The system as recited in claim 1, wherein the high availability
2 mechanism provides detection of, handling of, and recovery from a failure of the
3 forwarding component; the high availability mechanism including capabilities for
4 detection of a failure at the forwarding component by at least one load-balancing-
5 aware switch, redirection of packets to at least one other forwarding component,
6 and rebuilding of lost routes with a distributed session tracking manager.

7
8 4. The system as recited in claim 1, wherein the high availability
9 mechanism provides detection of, handling of, and recovery from a failure of the
10 classifying component; the high availability mechanism including capabilities for
11 detection of a failure at the classifying component by at least one forwarding
12 component, redirection of packets to at least one other classifying component, and
13 rebuilding of lost session information with a distributed session tracking manager.

14
15 5. The system as recited in claim 1, wherein the high availability
16 mechanism provides detection of, handling of, and recovery from a failure of the
17 session tracking component; the high availability mechanism including
18 capabilities for detection of a failure at the session tracking component by at least
19 one forwarding component and/or classifying component and for distributed and
20 redundant storage of session information.

1 6. The system as recited in claim 1, wherein the high availability
2 mechanism provides detection of, handling of, and recovery from a failure of the
3 health and load handling component; the high availability mechanism including
4 capabilities for detection of a failure at the health and load handling component by
5 at least one classifying component and for rebuilding of a cache of health and load
6 information using a message protocol.

7
8 7. The system as recited in claim 1, wherein the high availability
9 mechanism provides detection of, handling of, and recovery from a failure of the
10 health and load handling component; the high availability mechanism including
11 capabilities for redundant storing of health and load information and for
12 authoritative storing of health and load information at hosts to which the health
13 and load information pertains.

14
15 8. The system as recited in claim 1, wherein the forwarding component,
16 the classifying component, the session tracking component, and the health and
17 load handling component are resident at and executing on at least two different
18 devices.

1 **9.** An arrangement for highly available network load balancing
2 infrastructure, the arrangement comprising:

3 a plurality of different means for load balancing network traffic;

4 detection means for detecting a failure of one or more of the plurality of
5 different means for load balancing;

6 handling means for handling the failure; and

7 recovery means for recovering from the failure.

8
9 **10.** The arrangement as recited in claim 9, wherein the plurality of
10 different means for load balancing includes at least one forwarder means for
11 forwarding packets.

12
13 **11.** The arrangement as recited in claim 9, wherein the plurality of
14 different means for load balancing includes at least one classifier means for
15 classifying packets.

16
17 **12.** The arrangement as recited in claim 9, wherein the plurality of
18 different means for load balancing includes at least one request router means for
19 routing packets on a request-level.

20
21 **13.** The arrangement as recited in claim 9, wherein the plurality of
22 different means for load balancing includes at least one session tracker means for
23 tracking sessions.

1 **14.** The arrangement as recited in claim 9, wherein the plurality of
2 different means for load balancing includes at least one health and load handler
3 means for handling health and load information.

4
5 **15.** The arrangement as recited in claim 9, wherein the arrangement
6 comprises at least one system.

7
8 **16.** The arrangement as recited in claim 9, wherein the arrangement
9 comprises one or more processor-accessible media.

10
11 **17.** A network load balancing system comprising:
12 a first device that includes forwarding functionality; and
13 a second device that includes classifying functionality, the classifying
14 functionality performing classifying for the forwarding functionality;
15 wherein hardware of the first device differs from hardware of the second
16 device.

17
18 **18.** The network load balancing system as recited in claim 17, wherein
19 the hardware of the first device is especially tuned for the forwarding
20 functionality, and the hardware of the second device is especially tuned for the
21 classifying functionality.

1 **19.** The network load balancing system as recited in claim 17, wherein
2 the hardware of the first device has a relatively greater ability to accommodate a
3 high packet flux, and the hardware of the second device has a relatively greater
4 ability to accommodate processing-intensive tasks.

5
6 **20.** The network load balancing system as recited in claim 17, wherein
7 the hardware of the first device comprises a router or switch, and the hardware of
8 the second device comprises a personal computer or server.

9
10 **21.** One or more processor-accessible media comprising processor-
11 executable instructions that, when executed, enable a system to implement a route
12 plumbing protocol between a classifying component and a forwarding component
13 of network load balancing infrastructure, the route plumbing protocol usable for
14 communicating route information between the classifying component and the
15 forwarding component.

16
17 **22.** The one or more processor-accessible media as recited in claim 21,
18 wherein the route plumbing protocol is adapted to enable the classifying
19 component to send to the forwarding component an add route instruction and a
20 delete route instruction; the add route instruction causing the forwarding
21 component to add a route for a connection to a stipulated host, the delete route
22 instruction causing the forwarding component to delete a previously-added route.

1 **23.** The one or more processor-accessible media as recited in claim 21,
2 wherein the route plumbing protocol is adapted to enable the classifying
3 component to send to the forwarding component an add route instruction and a
4 delete route instruction from a first device including the classifying component to
5 a second device including the forwarding component.

6
7 **24.** A method comprising:
8 receiving a packet requesting a new connection at a forwarding component;
9 sending the packet from the forwarding component to a classifying
10 component;
11 selecting, by the classifying component, a route for the new connection;
12 and
13 plumbing, by the classifying component, the route for the new connection
14 by causing a new entry to be added in a local routing table of the forwarding
15 component.

16
17 **25.** The method as recited in claim 24, wherein the sending comprises:
18 sending the packet from a first device including the forwarding
19 component to a second device including the classifying component.

20
21 **26.** The method as recited in claim 24, wherein the selecting comprises:
22 selecting the route for the new connection responsive to a
23 preexisting session and based on a session identifier in the packet.
24
25

1 **27.** The method as recited in claim 24, wherein the selecting comprises:
2 selecting the route for the new connection responsive to health and
3 load information for a plurality of hosts.

4
5 **28.** The method as recited in claim 24, wherein the plumbing comprises:
6 sending an add route instruction from the classifying component to
7 the forwarding component.

8
9 **29.** The method as recited in claim 24, further comprising:
10 checking, by the forwarding component, the local routing table of the
11 forwarding component to determine that the packet is not for an existing
12 connection.

13
14 **30.** The method as recited in claim 24, further comprising:
15 adding, by the classifying component, a session information entry for the
16 route for the new connection in a distributed session-tracking table.

1 **31.** The method as recited in claim 24, further comprising:
2 adding, by the classifying component, a session information entry for the
3 route for the new connection in a distributed and replicated session-tracking table;
4 the new connection corresponding to a transmission control protocol/internet
5 protocol (TCP/IP) connection; the distributed and replicated session-tracking table
6 including at least one session information entry that corresponds to a higher-level
7 session in addition to the session information entry for the route for the new
8 connection.

9
10 **32.** The method as recited in claim 31, wherein the higher-level session
11 comprises a secure sockets layer (SSL) session.

12
13 **33.** A method comprising:
14 receiving a packet requesting a new session at a forwarding component;
15 sending the packet from the forwarding component to a classifying
16 component;
17 determining, by the classifying component, that the requested new session
18 fails to meet at least one firewall inspection policy; and
19 plumbing, by the classifying component, a route for the requested new
20 session that results in packets for the requested new session being dropped.

21
22 **34.** The method as recited in claim 33, wherein the requested new
23 session comprises a requested new connection.
24
25

1 **35.** One or more processor-accessible media comprising processor-
2 executable instructions that, when executed, direct a system to perform actions
3 comprising:

4 receiving a token allotment at traffic routing functionality from health and
5 load functionality, the token allotment having a first plurality of tokens for a first
6 destination and a second plurality of tokens for a second destination;

7 consuming, by the traffic routing functionality, a token of the first plurality
8 of tokens when selecting the first destination for a connection request; and

9 consuming, by the traffic routing functionality, a token of the second
10 plurality of tokens when selecting the second destination for a connection request.

11
12 **36.** The one or more processor-accessible media as recited in claim 35,
13 comprising the processor-executable instructions that, when executed, direct the
14 system to perform further actions comprising:

15 accumulating, by the health and load functionality, health and load
16 information from a plurality of hosts, the plurality of hosts including the first
17 destination and the second destination; and

18 determining, by the health and load functionality, the token allotment
19 responsive to the health and load information.
20
21
22
23
24
25

1 **37.** The one or more processor-accessible media as recited in claim 35,
2 comprising the processor-executable instructions that, when executed, direct the
3 system to perform a further action comprising:

4 selecting, by the traffic routing functionality, the token from between the
5 first plurality of tokens and the second plurality of tokens for the consuming using
6 a round-robin approach, a linear approach, or a current greater/greatest number of
7 tokens approach.

8
9 **38.** The one or more processor-accessible media as recited in claim 35,
10 wherein the first destination corresponds to a first application endpoint and the
11 second destination corresponds to a second application endpoint.

12
13 **39.** The one or more processor-accessible media as recited in claim 38,
14 wherein the first application endpoint and the second application endpoint
15 correspond to a same application type.

16
17 **40.** The one or more processor-accessible media as recited in claim 35,
18 wherein the traffic routing functionality comprises at least one of classifying
19 functionality and request routing functionality.

1 **41.** The one or more processor-accessible media as recited in claim 35,
2 comprising the processor-executable instructions that, when executed, direct the
3 system to perform a further action comprising:

4 receiving another token allotment at the traffic routing functionality from
5 the health and load functionality, the other token allotment having a third plurality
6 of tokens for a third destination and a fourth plurality of tokens for a fourth
7 destination.

8
9 **42.** The one or more processor-accessible media as recited in claim 41,
10 wherein the first destination and the second destination correspond to one
11 application type, and the third destination and the fourth destination correspond to
12 another application type.

1 **43.** One or more processor-accessible media comprising processor-
2 executable instructions that, when executed, direct a system to perform actions
3 comprising:

4 receiving, by a classifying component, a packet having a session reference
5 from a forwarding component;

6 searching, at least partly by the classifying component, a distributed session
7 tracking table to find a session information entry having a session identifier that
8 matches the session reference;

9 extracting, by the classifying component, a host identifier from the session
10 information entry having the session identifier that matches the session reference;
11 and

12 plumbing, by the classifying component, a route in a local routing table of
13 the forwarding component, the route for a connection that is associated with the
14 packet.

15
16 **44.** The one or more processor-accessible media as recited in claim 43,
17 comprising the processor-executable instructions that, when executed, direct the
18 system to perform a further action comprising:

19 forwarding, by the forwarding component, the packet to a host identified by
20 the host identifier.

1 **45.** The one or more processor-accessible media as recited in claim 43,
2 comprising the processor-executable instructions that, when executed, direct the
3 system to perform a further action comprising:

4 forwarding, by the classifying component, the packet to a host identified by
5 the host identifier.

6
7 **46.** The one or more processor-accessible media as recited in claim 43,
8 wherein the action of searching comprises an action of:

9 sending a query session information entry function call to at least
10 one other component.

11
12 **47.** The one or more processor-accessible media as recited in claim 43,
13 wherein the action of plumbing comprises an action of:

14 adding, as caused by the classifying component, a connection
15 identifier-host identifier pair to the local routing table of the forwarding
16 component.

17
18 **48.** The one or more processor-accessible media as recited in claim 47,
19 wherein the connection identifier of the connection identifier-host identifier pair
20 comprises a transmission control protocol (TCP) 4-tuple such that the connection
21 that is associated with the packet comprises a TCP connection that is treated as a
22 session by at least one session tracking component.

1 **49.** The one or more processor-accessible media as recited in claim 47,
2 comprising the processor-executable instructions that, when executed, direct the
3 system to perform a further action comprising:

4 adding, by the classifying component, the connection identifier-host
5 identifier pair to the distributed session tracking table.

6
7 **50.** One or more processor-accessible media comprising processor-
8 executable instructions for load balancing infrastructure that, when executed,
9 enable a system to perform actions comprising:

10 establishing a first connection with a client;

11 receiving a first request from the client via the first connection;

12 determining, responsive to session information and/or health and load
13 information, that the first request is to be routed to a first host via a second
14 connection;

15 receiving a second request from the client via the first connection; and

16 determining, responsive to the session information and/or the health and
17 load information, that the second request is to be routed to a second host via a
18 third connection.

19
20 **51.** The one or more processor-accessible media as recited in claim 50,
21 comprising the processor-executable instructions that, when executed, enable the
22 system to perform further actions comprising:

23 establishing the second connection to the first host; and

24 establishing the third connection to the second host.
25

1 **52.** The one or more processor-accessible media as recited in claim 51,
2 wherein the action of establishing the second connection can be performed prior to
3 the action of receiving a first request.
4

5 **53.** The one or more processor-accessible media as recited in claim 51,
6 wherein the actions of establishing comprise actions of establishing a transmission
7 control protocol (TCP) connection.
8

9 **54.** The one or more processor-accessible media as recited in claim 50,
10 comprising the processor-executable instructions that, when executed, enable the
11 system to perform further actions comprising:

12 modifying the first request; and

13 routing the modified first request to the first host via the second connection.
14

15 **55.** The one or more processor-accessible media as recited in claim 54,
16 wherein the action of modifying comprises at least one action of:

17 decrypting the first request; and

18 aggregating the first request with one or more other requests.
19
20
21
22
23
24
25

1 **56.** A system comprising:
2 a first classifying component that classifies packets;
3 a second classifying component that classifies packets; and
4 a forwarding component that forwards packets, the forwarding component
5 assigned to the first classifying component to attain classification services
6 therefrom;

7 wherein the system is adapted to change the forwarding component to be
8 assigned to the second classifying component when a failure of the first classifying
9 component is detected.

10
11 **57.** The system as recited in claim 56, wherein the forwarding
12 component is capable of detecting the failure of the first classifying component.

13
14 **58.** The system as recited in claim 56, wherein the forwarding
15 component is capable of reassigning itself to the second classifying component
16 when the failure of the first classifying component is detected.

17
18 **59.** The system as recited in claim 56, wherein the first classifying
19 component and the second classifying component classify packets responsive to
20 session information and/or health and load information.

21
22 **60.** The system as recited in claim 56, wherein the system comprises at
23 least one of a network load balancer and a firewall.

1 **61.** One or more processor-accessible media comprising processor-
2 executable instructions that, when executed, enable a system to perform actions
3 comprising:

4 receiving a first packet for a connection at first forwarding functionality;
5 plumbing a route for the connection at the first forwarding functionality;
6 receiving a second packet for the connection at second forwarding
7 functionality; and
8 plumbing the route for the connection at the second forwarding
9 functionality using a distributed session tracking table.

10
11 **62.** The one or more processor-accessible media as recited in claim 61,
12 wherein the actions of plumbing a route for the connection at the first forwarding
13 functionality and plumbing the route for the connection at the second forwarding
14 functionality are performed by a single classifying functionality.

15
16 **63.** The one or more processor-accessible media as recited in claim 61,
17 wherein the action of plumbing a route for the connection at the first forwarding
18 functionality is performed by first classifying functionality, and the action of
19 plumbing the route for the connection at the second forwarding functionality is
20 performed by second classifying functionality.

1 **64.** The one or more processor-accessible media as recited in claim 61,
2 comprising the processor-executable instructions that, when executed, enable the
3 system to perform a further action comprising:

4 experiencing a failure at the first forwarding functionality prior to the action
5 of receiving a second packet for the connection at second forwarding functionality.
6

7 **65.** The one or more processor-accessible media as recited in claim 61,
8 wherein the action of plumbing the route for the connection at the second
9 forwarding functionality using a distributed session tracking table comprises an
10 action of:

11 plumbing the route for the connection at the second forwarding
12 functionality using the distributed session tracking table that is replicated to
13 at least one level.
14

15 **66.** The one or more processor-accessible media as recited in claim 61,
16 wherein the action of plumbing the route for the connection at the second
17 forwarding functionality using a distributed session tracking table comprises an
18 action of:

19 making a query connection function call to a distributed session
20 tracking manager that manages the distributed session tracking table.
21
22
23
24
25

1 **67.** One or more processor-accessible media comprising processor-
2 executable instructions that, when executed, direct a system to perform actions
3 comprising:

4 receiving a message having a content indicator for a health and/or load
5 table of a host;

6 comparing the content indicator for the health and/or load table of the host
7 to at least part of local cache contents; and

8 if the at least part of local cache contents are not equivalent to the content
9 indicator for the health and/or load table of the host, sending a message requesting
10 content of the health and/or load table of the host.

11
12 **68.** The one or more processor-accessible media as recited in claim 67,
13 comprising the processor-executable instructions that, when executed, direct the
14 system to perform a further action comprising:

15 receiving a message having the requested content of the health and/or load
16 table of the host.

17
18 **69.** The one or more processor-accessible media as recited in claim 68,
19 wherein the message having the requested content of the health and/or load table
20 of the host comprises a send table snapshot message.

21
22 **70.** The one or more processor-accessible media as recited in claim 67,
23 wherein the actions of receiving and sending comprise the respective actions of
24 receiving from the host and sending towards the host.

25

1 **71.** The one or more processor-accessible media as recited in claim 67,
2 wherein the content indicator for the health and/or load table of the host comprises
3 error checking data.

4
5 **72.** The one or more processor-accessible media as recited in claim 67,
6 wherein the content indicator for the health and/or load table of the host
7 corresponds to at least one of (i) an entirety of the health and/or load table of the
8 host and (ii) a portion of the health and/or load table of the host.

9
10 **73.** The one or more processor-accessible media as recited in claim 67,
11 wherein the message having the content indicator for the health and/or load table
12 of the host comprises a heartbeat message.

13
14 **74.** The one or more processor-accessible media as recited in claim 67,
15 wherein the message requesting content of the health and/or load table of the host
16 comprises a get table snapshot message.